arXiv:1510.04361v2 [math.NA] 27 Jul 2016

Global sensitivity metrics from active subspaces

Paul G. Constantine[a,*], Paul Diaz[b]

[a]*Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO*

[b]*Department of Aerospace Engineering, University of Colorado Boulder, Boulder, CO*

Abstract

Predictions from science and engineering models depend on several input parameters. Global sensitivity analysis quantifies the importance of each input parameter, which can lead to insight into the model and reduced computational cost; commonly used sensitivity metrics include Sobol' total sensitivity indices and derivative-based global sensitivity measures. *Active subspaces* are an emerging set of tools for identifying important directions in a model's input parameter space; these directions can be exploited to reduce the model's dimension enabling otherwise infeasible parameter studies. In this paper, we develop global sensitivity metrics called *activity scores* from the active subspace, which yield insight into the important model parameters. We mathematically relate the activity scores to established sensitivity metrics, and we discuss computational methods to estimate the activity scores. We show two numerical examples with algebraic functions taken from simplified engineering models. For each model, we analyze the active subspace and discuss how to exploit the low-dimensional structure. We then show that input rankings produced by the activity scores are consistent with rankings produced by the standard metrics.

Keywords: sensitivity analysis, active subspaces, activity score

*Corresponding author

Email address: paul.constantine@mines.edu (Paul G. Constantine)

URL: http://inside.mines.edu/~pconstan (Paul G. Constantine)

1. Introduction

Complex science and engineering models depend on several input parameters. The complexity often creates uncertainty about which parameters are most important for describing physical processes and model predictions. Sensitivity analysis seeks to identify the most important parameters for particular outputs of interest; effectively identifying important parameters leads to insight into the model [1]. Sensitivity analysis may enable the scientist to ignore less important parameters, thus reducing the cost of parameter studies—e.g., optimization or uncertainty quantification—that are essential for confident predictions. Sensitivity metrics are broadly classified as *local*, which measure the model's response to small perturbations around a nominal parameter value, and *global*, which assess the importance of each variable over a range of parameters; we are primarily concerned with global sensitivity metrics. The *total sensitivity indices* of Sobol' are some of the most widely used metrics for global sensitivity analysis [2]. Other common metrics average local gradients to build global metrics [3]. The appropriate choice of metric depends on the application.

Using sensitivity metrics to ignore less important parameters is a form of dimension reduction in the model's input parameter space. This type of reduction is limited to the coordinates induced by the model parameters. Our recent work in *active subspaces* generalizes such parameter space dimension reduction to directions in the space of model inputs [4, 5]. A model's output might depend strongly on each coordinate but only through a limited number of directions. Each direction corresponds to a linear combination of the input parameters that we call an *active variable*. Our work has developed methods for identifying an active subspace [6] and exploiting it, if present, to reduce the dimension of subsequent parameter studies [4]. We have applied this approach to aerospace shape optimization [7], automotive design [8], uncertainty quantification for multiphysics scramjet models [9], sensitivity analysis in integrated hydrologic models [10], global spatial sensitivity in subsurface flow [11], dimension reduction in single-diode solar cell models [12], uncertainty quantification for combustion models [13], time-dependent sensitivity analysis for lithium ion battery models [14], and sensitivity analysis for Ebola virus population dynamics [15]. The number of input parameters in these models varies between 5 and 100. The active subspace quantifies low-dimensional structure in the input/output map of a given model. The aforementioned list of science and engineering applications

demonstrates that such structures are present in many real world models; moreover we have shown that these low-dimensional structures are generically present in any physical system with units [16]. If an active subspace is sufficiently low dimensional in a given model, then the researcher may exploit the active subspace to enable otherwise infeasible parameter studies. The way in which she exploits the structure depends on the study. For example, to find the range of possible outputs over the parameter space (i.e., minimization and maximization), she may exploit a global, monotonic trend in the relationship between the active variables and the output quantity of interest; we apply this approach in [9] to find the safe operating in multiphysics scramjet model. Alternatively, if she wishes to construct a cheap response surface that mimics the model's input/output relationship, she may train the surface on only the active subspace coordinates; see [17] for details.

One difficulty in subspace-based dimension reduction is how to draw insight into the original model parameters once an active subspace has been identified. For example, how might a climate scientist use the fact that $3 \times$ (cloud parameter) $+ 2 \times$ (radiation parameter) is an important linear combination in a climate model? In this paper, we propose global sensitivity metrics based on active subspaces; our main goal is to demonstrate that the proposed metrics function like standard sensitivity metrics—even though the proposed metrics' derivation and motivation are quite different. Toward this goal, we mathematically relate the proposed metrics to existing sensitivity metrics: Sobol' total sensitivity indices and derivative-based global sensitivity measures. We also discuss two numerical procedures for estimating the proposed metrics: tensor product Gauss quadrature and Monte Carlo. For the Monte Carlo method, we propose bootstrap-based standard error estimates. We use two simplified engineering models to compare the proposed metrics to existing metrics: (i) a 7-parameter algebraic model of piston cycle time and (ii) a 6-parameter algebraic model of a transformerless circuit voltage. The computationally inexpensive models with relatively few parameters allow us to perform a thorough numerical error study, where we compare the Monte Carlo-based estimates to quadrature-based reference values. We also numerically compare the proposed metrics to the coefficients of a least-squares fit linear model; the related standard regression coefficients are common sensitivity metrics. We show that the rankings produced by all metrics are similar for both models; this result supports the claim that the active subspace-based sensitivity metrics are comparable to the standard metrics. It is not our goal to demonstrate that the active subspace-based

metrics are superior to existing metrics—either in characterizing sensitivity for a given model or in computational cost for estimating the metrics. In fact, when derivatives are available, the derivative-based global sensitivity measures use a comparable number of gradient evaluations, and they do not require an eigendecomposition.

The remainder of the paper proceeds as follows. Section 2 reviews three common sensitivity metrics: Sobol' total sensitivity indices, derivative-based global sensitivity measures, and linear model coefficients. This section sets up notation for the remainder of the paper. Section 3 reviews active subspaces. Section 4 mathematically defines the proposed sensitivity metrics derived from the active subspace and relates the proposed metrics to the total sensitivity indices and the derivative-based global sensitivity measures. Section 5 shows the results of the numerical experiments. And Section 6 summarizes the results, discusses limitations of the proposed metrics, and offers directions for future research.

1.1. Setup and notation

To study the sensitivity metrics derived from active subspaces, we consider a differentiable and square-integrable function $f = f(\mathbf{x})$ defined on an m-dimensional hypercube $\mathbf{x} = [x_1, \ldots, x_m]^T \in [-1, 1]^m$. Define the uniform probability density function $\rho(\mathbf{x}) = 2^{-m}$ for $\mathbf{x} \in [-1, 1]^m$ and zero elsewhere. The mean and variance of f are

$$\mathbb{E}[f] = \int f(\mathbf{x})\,\rho(\mathbf{x})\,d\mathbf{x}, \qquad \mathrm{Var}[f] = \int \big(f(\mathbf{x}) - \mathbb{E}[f]\big)^2 \rho(\mathbf{x})\,d\mathbf{x}. \quad (1)$$

We emphasize that f is a deterministic function; its mean and variance do not represent any probabilistic information. Denote the gradient of f by $\nabla f(\mathbf{x}) \in \mathbb{R}^m$ with partial derivatives $\frac{\partial f}{\partial x_i}(\mathbf{x})$, and assume the partial derivatives are square-integrable over the hypercube domain. All vector norms are the standard Euclidean norm, and matrix norms are the matrix 2-norm (i.e., the spectral norm).

This set up is an idealized input/output map. To formally apply the following methods and results to a real application, the researcher must (i) identify a scalar-valued quantity of interest that satisfies the assumptions (e.g., square-integrable and differentiable), (ii) identify a vector of continuous input parameters, and (iii) transform the input parameter space to the hypercube with the uniform measure.

2. Global sensitivity metrics

We briefly review three global sensitivity metrics used in practice: variance-based total sensitivity indices, derivative-based global sensitivity measures, and least-squares-fit linear model coefficients. In each case, we (i) define the metric, (ii) offer a method for high accuracy computation with numerical quadrature, (iii) detail a Monte Carlo method for low accuracy estimation, and (iv) mention practical error estimates for the Monte Carlo estimates. We use lower case Greek letters to denote the various sensitivity metrics, where each symbol's subscript identifies its corresponding variable (i.e., component of the vector of inputs $\mathbf{x}$). The book by Saltelli, et al. [1] provides an overview of global sensitivity analysis and its metrics including common metrics we do not review, such as correlation coefficients.

2.1. Sobol' total sensitivity indices

The total sensitivity indices were introduced by Sobol' [2]; several papers describe these metrics, propose extensions [18], and develop efficient computational procedures [19, 20, 21, 22, 23, 24]. The total sensitivity indices are derived from a particular decomposition of the square-integrable $f(\mathbf{x})$ called the functional ANOVA decomposition or the variance-based decomposition; we use the concise notation from Liu and Owen [19] to describe the decomposition and the derived sensitivity indices. Let u be a subset of the integers $\{1, \ldots, m\}$ with cardinality $|u|$. Define $\mathbf{x}^u \in [-1, 1]^{|u|}$. The subset u_c is the complement of u in $\{1, \ldots, m\}$ with cardinality $|u_c| = m - |u|$. The ANOVA decomposition can be written

$$f(\mathbf{x}) \;=\; \sum_{u \subseteq \{1,\ldots,m\}} f_u(\mathbf{x}), \tag{2}$$

where f_u depends on $\mathbf{x}$ only through $\mathbf{x}^u$. Each function f_u is defined as

$$f_u(\mathbf{x}) \;=\; 2^{-|u_c|} \int_{[-1,1]^{|u_c|}} f(\mathbf{x})\, d\mathbf{x}^{u_c} - \sum_{v \subsetneq u} f_v(\mathbf{x}), \tag{3}$$

where $f_\emptyset = \mathbb{E}[f]$. In the standard construction, the functions f_u are mutually orthogonal. The variance of f can be decomposed as

$$\mathrm{Var}[f] \;=\; \sum_{u \subseteq \{1,\ldots,m\}} \mathrm{Var}[f_u]\,. \tag{4}$$

Define the set $\mathcal{S}_i$ to be the set of subsets of $\{1, \ldots, m\}$ containing the index i. The total sensitivity index τ_i for the ith variable x_i can be written as sums of subsets of the partial variances $\mathrm{Var}[f_u]$ divided by the total variance,

$$\tau_i \;=\; \frac{\sum_{u \subseteq \mathcal{S}_i} \mathrm{Var}[f_u]}{\mathrm{Var}[f]}. \tag{5}$$

Some authors use the notation S_{T_i} [1] or S_i^{tot} [2] to denote the total sensitivity indices. We use the lower case Greek letter τ_i for consistency. Other sensitivity metrics can be defined from the variances $\mathrm{Var}[f_u]$; we focus exclusively on τ_i from (5) to compare with the metrics derived from active subspaces; see Section 4.

When the dimension of $\mathbf{x}$ is sufficiently low, $f(\mathbf{x})$ is sufficiently smooth, and $f(\mathbf{x})$ can be evaluated quickly, one can accurately estimate τ_i using a high order numerical quadrature rule, such as a tensor product Gauss-Legendre quadrature rule. This approach is equivalent to using the Legendre series' coefficients as in [21, 22], which we use to compute reference values for the numerical examples in Section 5.

If $f(\mathbf{x})$ does not satisfy these assumptions (e.g., if the dimension of $\mathbf{x}$ is too large to permit tensor product quadrature), one can use a less accurate Monte Carlo method as in [1, Section 4.6]. Several alternative Monte Carlo methods have been proposed; see [23, 24] for a detailed discussion. We review Jansen's method [25], which is the method we use for numerical experiments, and a bootstrap standard error. First draw $2M$ independent samples $\mathbf{x}_i$ according to the uniform density $\rho(\mathbf{x})$ on $[-1, 1]^m$. Arrange these samples into two $M \times m$ matrices, $\boldsymbol{A}$ and $\boldsymbol{B}$, defined as

$$\boldsymbol{A} = \begin{bmatrix} \mathbf{x}_1^T \\ \vdots \\ \mathbf{x}_M^T \end{bmatrix}, \qquad \boldsymbol{B} = \begin{bmatrix} \mathbf{x}_{M+1}^T \\ \vdots \\ \mathbf{x}_{2M}^T \end{bmatrix}. \tag{6}$$

Define the $M \times m$ matrix $\boldsymbol{B}_i$ to be the matrix $\boldsymbol{B}$ with the ith column replaced by the ith column of $\boldsymbol{A}$. Define the M-vector $\mathbf{f}_{\boldsymbol{A}}$ as

$$\mathbf{f}_{\boldsymbol{A}} \;=\; \begin{bmatrix} f(\mathbf{x}_1) \\ \vdots \\ f(\mathbf{x}_M) \end{bmatrix}. \tag{7}$$

In words, the vector $\mathbf{f}_{\boldsymbol{A}}$ contains evaluations of the function f at each row of the matrix $\boldsymbol{A}$. Similarly, define the M-vectors $\mathbf{f}_{\boldsymbol{B}_i}$ whose elements contain

evaluations of f at the rows of their respective subscripted matrices. These calculations use $M(1+m)$ unique evaluations of f. The total sensitivity indices may be estimated with Jansen's formula [25, 23],

$$\tau_i \;\approx\; \hat{\tau}_i \;=\; \frac{\|f_{\boldsymbol{A}} - f_{\boldsymbol{B}_i}\|^2}{2\,M\,\hat{\sigma}_f^2}, \tag{8}$$

where $\hat{\sigma}_f^2$ is the sample variance of the M evaluations of f in $\mathbf{f}_{\boldsymbol{A}}$, which approximates the true variance of f.

Under certain assumptions on the partial variances $\mathrm{Var}\,[f_u]$ from (4), one can derive the standard error of the Monte Carlo estimator from a chi-squared distribution [25]. However, these assumptions are not satisfied for a general nonlinear $f(\mathbf{x})$. To estimate the standard error, we employ a nonparametric bootstrap [26]. We denote the bootstrap standard error for the Monte Carlo estimate $\hat{\tau}_i$ by $\widehat{se}_{\tau_i}$. And we numerically study the bootstrap standard error as a function of the number of samples in Section 5.

2.2. Derivative-based global sensitivity measures

A natural notion of sensitivity is how a model output responds to small changes in its inputs, i.e., a derivative. However, derivatives only provide information at a single parameter value. Morris developed sensitivity metrics based on sampling and averaging coarse finite difference approximations to derivatives [27]. Kucherenko, et al. [3] extended this idea to continuous derivatives and integration, where the the averages are estimated with Monte Carlo methods. We focus on one particular derivative-based sensitivity metric studied by Sobol' and Kucherenko [28]:

$$\nu_i \;=\; \int \left(\frac{\partial f}{\partial x_i}(\mathbf{x})\right)^2 \rho(\mathbf{x})\,d\mathbf{x}, \qquad i = 1, \ldots, m. \tag{9}$$

This metric averages local sensitivity information to make it a global metric. We discuss this metric's connection to active subspaces in Section 4.

If f's derivatives are sufficiently smooth, and if the dimension of $\mathbf{x}$ is low, then one can accurately estimate ν_i using high order numerical quadrature, such as tensor-product Gauss-Legendre quadrature; we use the quadrature approach for reference values in the numerical examples in section 5.

The Monte Carlo estimate of ν_i using M samples is

$$\nu_i \;\approx\; \hat{\nu}_i = \frac{1}{M}\sum_{j=1}^{M}\left(\frac{\partial f}{\partial x_i}(\mathbf{x}_j)\right)^2, \tag{10}$$

where the $\mathbf{x}_j$'s are drawn independently according to $\rho(\mathbf{x})$. The standard error se_{ν_i} is

$$se_{\nu_i} \;=\; \left[\frac{1}{M-1}\sum_{j=1}^{M}\left(\left(\frac{\partial f}{\partial x_i}(\mathbf{x}_j)\right)^2 - \hat{\nu}_i\right)^2\right]^{1/2}. \tag{11}$$

In section 5, we study the standard error of the estimator by investigating se_{ν_i} for different values of M.

2.3. Linear model coefficients

A simple and easily computable sensitivity metric can be derived from the coefficients of a least-squares fit linear approximation to $f(\mathbf{x})$. Such a metric is valid when f is smooth and monotonic along each component of $\mathbf{x}$. Saltelli, et al. [1, Section 1.2.5] discuss a similar metric in connection with a linear regression model. Since $f(\mathbf{x})$ is deterministic, the statistical interpretation of the regression coefficients is not valid. Nevertheless, we can devise a Monte Carlo-based least-squares procedure to estimate the coefficients of a linear approximation as follows. Draw $\mathbf{x}_j$ independently according to $\rho(\mathbf{x})$ for $j = 1, \ldots, M$, and compute $f_j = f(\mathbf{x}_j)$. Next, use least-squares to fit the intercept b_0 and coefficients $\mathbf{b} = [b_1, \ldots, b_m]^T$ of a linear approximation,

$$f_j \;\approx\; b_0 + \mathbf{b}^T\mathbf{x}_j, \qquad j = 1, \ldots, M. \tag{12}$$

Scaling the coefficients b_i as follows produces useful sensitivity metrics called the *standardized regression coefficients*; see [1, Section 1.2.5]. Define β_i as

$$\beta_i \;=\; \frac{b_i}{\sqrt{3}\,\hat{\sigma}_f}, \tag{13}$$

where $\hat{\sigma}_f$ is the standard deviation of the set of samples $\{f_j\}$. The $1/\sqrt{3}$ comes from scaling the coefficients by the variance of the uniformly distributed random variable x_i with support $[-1, 1]$, which is consistent with the problem set up in Section 1.1. If $f(\mathbf{x})$ is monotonic and sufficiently smooth, then β_i gives a rough indication of how f changes in response to a change in x_i over the space of inputs. In contrast to the total sensitivity index τ_i from (8) and the derivative-based measures ν_i from (9), the β_i's are signed, so they indicate if f will increase or decrease given a positive change in x_i.

We can compute reference values for β_i by noting the relationship between the m-dimensional Legendre series truncated to degree 1 (i.e., only a constant and m linear terms) and the least-squares linear approximation (12); the Legendre series is often referred to as the *generalized polynomial chaos* associated with the uniform density [21, 29]. In particular, the degree-1 Legendre approximation is the best linear approximation in the L_2 norm. Its coefficients admit a simple integral representation by virtue of the Legendre polynomials' orthogonality. Thus, we can estimate the coefficients with high order Gauss quadrature [30] and scale them by a quadrature-based estimate of the variance to get the coefficients of a linear monomial approximation as in (12).

When reference values are too expensive to compute, we can estimate the standard error in β_i with a bootstrap standard error, which we denote $\widehat{se}_{\beta_i}$. See [26, Section 9.5] for a detailed description of using the bootstrap for linear regression coefficients. Note that randomness in the estimates is a result of the Monte Carlo method; there is no randomness in the evaluation of $f(\mathbf{x})$. In section 5, we numerically study the decay in the bootstrap standard error for different values of M.

3. Active subspaces

The active subspace [5] is defined by the eigenvectors of the following $m \times m$ symmetric, positive semidefinite matrix,

$$\boldsymbol{C} \;=\; \int \nabla f(\mathbf{x})\, \nabla f(\mathbf{x})^T\, \rho(\mathbf{x})\, d\mathbf{x} \;=\; \boldsymbol{W} \Lambda \boldsymbol{W}^T, \tag{14}$$

where $\boldsymbol{W} = [\mathbf{w}_1, \ldots, \mathbf{w}_m]$ is the orthogonal matrix of eigenvectors, and $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_m)$ is the diagonal matrix of eigenvalues in decreasing order. The eigenpairs are functionals, i.e., properties of the given f, similar to the total sensitivity indices or derivative-based measures. The eigenpairs satisfy

$$\lambda_i \;=\; \mathbf{w}_i^T \boldsymbol{C} \mathbf{w}_i \;=\; \int \left(\nabla f(\mathbf{x})^T \mathbf{w}_i\right)^2 \rho(\mathbf{x})\, d\mathbf{x}, \tag{15}$$

which implies that $\lambda_i = 0$ if and only if f is constant along the direction $\mathbf{w}_i$. If such structure is present in the given f, then the structure can be exploited to reduce the parameter space dimension for parameter studies. Suppose $\lambda_n > \lambda_{n+1}$ for some $n < m$. Then we can partition the eigenpairs,

$$\Lambda = \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix}, \qquad \boldsymbol{W} = \begin{bmatrix} \boldsymbol{W}_1 & \boldsymbol{W}_2 \end{bmatrix}, \tag{16}$$

where Λ_1 contains the first n eigenvalues, and $\boldsymbol{W}_1$ is the $m \times n$ matrix containing the first n eigenvectors. The *active subspace* of dimension n is the column span of $\boldsymbol{W}_1$. Not all square-integrable and differentiable functions admit an active subspace. For example, for the function $f(\mathbf{x}) = \|\mathbf{x}\|^2$, which is radially symmetric, and a uniform density $\rho(\mathbf{x})$, the eigenvalues are are all equal.

If $\lambda_{n+1}, \ldots, \lambda_m$ are sufficiently small, then a reasonable low-dimensional approximation of f is

$$f(\mathbf{x}) \;\approx\; g(\boldsymbol{W}_1^T \mathbf{x}), \tag{17}$$

where $g : \mathbb{R}^n \to \mathbb{R}$. In [4], we propose a particular form for g based on a conditional average, and we bound its mean-squared error in terms of the eigenvalues $\lambda_{n+1}, \ldots, \lambda_m$.

To identify an active subspace, we must estimate the eigenpairs Λ, $\boldsymbol{W}$ and examine the estimated eigenvalues. If the dimension m is sufficiently small, then one can use high order numerical integration (e.g., tensor product Gauss-Legendre quadrature) to estimate $\boldsymbol{C}$ in (14). In [6], we propose and analyze the following Monte Carlo method for estimating $\boldsymbol{C}$'s eigenpairs. For $i = 1, \ldots, M$, draw $\mathbf{x}_i$ independently according to $\rho(\mathbf{x})$. Compute $\nabla f_i = \nabla f(\mathbf{x}_i)$. Then estimate

$$\boldsymbol{C} \;\approx\; \hat{\boldsymbol{C}} \;=\; \frac{1}{M} \sum_{i=1}^{M} \nabla f_i \, \nabla f_i^T \;=\; \hat{\boldsymbol{W}} \hat{\Lambda} \hat{\boldsymbol{W}}^T. \tag{18}$$

In [6], we study how large M must be so that (i) the estimated eigenvalues $\hat{\Lambda}$ are close to the true eigenvalues Λ and (ii) the subspace distance ε defined as

$$\varepsilon \;=\; \|\boldsymbol{W}_1 \boldsymbol{W}_1^T - \hat{\boldsymbol{W}}_1 \hat{\boldsymbol{W}}_1^T\|, \tag{19}$$

where $\hat{\boldsymbol{W}}_1$ contains the first n columns of $\hat{\boldsymbol{W}}$, is well-behaved. The main result from [6] is Corollary 3.7, which shows that

$$\varepsilon \;\leq\; \frac{4 \, \lambda_1 \, \delta}{\lambda_n - \lambda_{n+1}}, \tag{20}$$

where δ is a user-specified tolerance on the relative eigenvalue error that determines the minimum number of required samples. Equation (20) shows that a large eigenvalue gap $\lambda_n - \lambda_{n+1}$ implies that the active subspace can be accurately estimated with Monte Carlo. Therefore, the practical heuristic is to choose the dimension n of the active subspace according to the largest gap in the eigenvalues.

4. Sensitivity metrics from active subspaces

We propose two sensitivity metrics derived from the eigenvectors $\boldsymbol{W}$ and eigenvalues Λ from (14). The first metric uses the components of the first eigenvector (i.e., the eigenvector associated with the largest eigenvalue). The second metric uses a linear combination of the squared eigenvector components weighted by the eigenvalues, which we call *activity scores*.

4.1. The first eigenvector

In [12, Section 4.4], we used the first eigenvector from the matrix $\hat{\boldsymbol{C}}$ in (18) to identify the most important parameters in a single-diode model of a photovoltaic solar cell. For this particular model, the magnitudes of the eigenvector components gave the same ranking as the total sensitivity indices (see section 2.1). However, these two metrics measure different characteristics of the function. The total sensitivity indices measure the proportion of the variance attributed to each parameter while the eigenvector identifies an important direction in the parameter space, where "importance" is measured as (15). In words, perturbing $\mathbf{x}$ along $\mathbf{w}_1$ changes $f(\mathbf{x})$ the most, on average.

Like the linear model coefficients in section 2.3, the eigenvector components are not all positive, and the relative difference in signs (since eigenvectors are only unique up to a sign) provides insight into the relationship between inputs and outputs. For example, if two eigenvector components are equal in magnitude but differ in sign, then we expect f to increase when one parameter is positively perturbed but decrease when the other is positively perturbed. This is similar to the gradient-based metrics described by Morris [27] and Kucherenko [3]. However, the trouble of averaging a gradient that might change signs does not effect the eigenvector. In the numerical examples in section 5, we compare the first eigenvector's components to other sensitivity metrics, but we do not provide a theoretical comparison. We use the bootstrap to estimate a standard error in the eigenvector components, though one must be careful to keep the signs consistent when computing the bootstrap replicates. Section 7.2 of [26] shows an example of using the bootstrap to estimate standard errors in eigenvector components. We denote the first eigenvector by $\mathbf{w}_1 = [w_{1,1}, \ldots, w_{m,1}]^T$ and the bootstrap standard error of its ith component by $\widehat{se}_{w_{i,1}}$.

4.2. Activity scores

We develop a more comprehensive and justifiable sensitivity metric from the eigenpairs in (14). We propose the following metric for global sensitivity

analysis. Define α_i as

$$\alpha_i \;=\; \alpha_i(n) \;=\; \sum_{j=1}^{n} \lambda_j \, w_{i,j}^2, \qquad i = 1, \ldots, m. \tag{21}$$

We call $\alpha_i(n)$ the *activity score* for the ith parameter, and we use these numbers to rank the importance of a model's inputs.

The active subspace's construction gives insight into the activity scores' interpretation. The eigenvector $\mathbf{w}_1$ identifies the most important direction in the parameter space in the following sense: perturbing $\mathbf{x}$ along $\mathbf{w}_1$ changes f more, on average, than perturbing $\mathbf{x}$ orthogonal to $\mathbf{w}_1$; see (15). The components of $\mathbf{w}_1$ measure the relative change in each component of $\mathbf{x}$ along this most important direction, so they impart significance to each component of $\mathbf{x}$. The second most important direction is the eigenvector $\mathbf{w}_2$, and the relative importance of $\mathbf{w}_2$ is measured by the difference between the eigenvalues λ_1 and λ_2. For example, if $\lambda_1 \gg \lambda_2$, then $f(\mathbf{x})$ has a one-dimensional active subspace, and the importance of $\mathbf{x}$'s components is captured in $\mathbf{w}_1$'s components. Therefore, to construct the global sensitivity metric, it is reasonable to scale each eigenvector by its corresponding eigenvalue. Squaring each eigenvector component in (21) removes information provided by the signs. But the resulting metric is much easier to compare to existing sensitivity metrics.

4.3. Comparison to existing metrics

We compare the activity scores to the derivative-based metrics and the total sensitivity indices reviewed in Section 2.

Theorem 4.1. *The activity scores are bounded above by the derivative-based metrics,*

$$\alpha_i(n) \;\leq\; \nu_i, \qquad i = 1, \ldots, m, \tag{22}$$

where ν_i is from (9)*. The inequality becomes an equality when $n = m$.*

Proof. Note that the derivative-based metrics are the diagonal elements of the matrix $\boldsymbol{C}$ in (14). Let $\boldsymbol{\nu} = [\nu_1, \ldots, \nu_m]^T$. Then

$$\boldsymbol{\nu} \;=\; \mathrm{diag}(\boldsymbol{C}) \;=\; \mathrm{diag}(\boldsymbol{W}\Lambda\boldsymbol{W}^T). \tag{23}$$

So

$$\nu_i \;=\; \sum_{j=1}^{m} \lambda_j \, w_{i,j}^2 \;=\; \alpha_i(m), \qquad i = 1, \ldots, m, \tag{24}$$

where $w_{i,j}$ is the (i,j) element of $\boldsymbol{W}$. This proves the equality statement. To see the inequality, note that $w_{i,j}^2 \geq 0$ and $\lambda_j \geq 0$, so

$$\alpha_i(n) \;=\; \sum_{j=1}^{n} \lambda_j \, w_{i,j}^2 \;\leq\; \sum_{j=1}^{m} \lambda_j \, w_{i,j}^2 \;=\; \nu_i, \tag{25}$$

as required.

Theorem 4.1 shows that the activity score can be interpreted as a truncation of the derivative-based metric ν_i from (9). In many practical situations, the activity scores and the derivative-based metrics give comparable rankings. However, it is possible to construct cases where the rankings differ. Consider the quadratic function $f(x_1, x_2) = (x_1^2 + x_2^2)/2$. For this case, the matrix $\boldsymbol{C}$ from (14) is

$$\boldsymbol{C} \;=\; \begin{bmatrix} 1/3 & 0 \\ 0 & 1/3 \end{bmatrix}. \tag{26}$$

The derivative-based metrics are the diagonal elements of $\boldsymbol{C}$—both $1/3$ implying equal importance in the two inputs. However, the activity scores with $n = 1$ are $[1/3, 0]$, which implies that the second variable is not important at all. This case illustrates that the activity scores are appropriate when there is a gap between the eigenvalues of $\boldsymbol{C}$ (i.e., f admits an active subspace), and n is chosen according to the gap. In the case of (26), there is no gap between the two eigenvalues, so it is not possible to distinguish a one-dimensional eigenspace that defines a one-dimensional active subspace.

Sobol' and Kucherenko connect the total sensitivity indices to the derivative-based metrics in [28]; we use [28, Theorem 2] to relate the total sensitivity indices to the activity scores.

Theorem 4.2. *The total sensitivity index τ_i from* (8) *is bounded by*

$$\tau_i \;\leq\; \frac{1}{4\pi^2 \, V} \, (\alpha_i(n) + \lambda_{n+1}), \qquad i = 1, \ldots, m, \tag{27}$$

where $V = \operatorname{Var}[f]$, $\alpha_i(n)$ is from (21)*, and λ_{n+1} is the $n+1$st eigenvalue in* (14)*.*

Proof. Theorem 2 from [28] shows

$$\tau_i \;\leq\; \frac{1}{4\pi^2 \, V} \, \nu_i, \tag{28}$$

for $i = 1, \ldots, m$. Note that an additional factor of $1/4$ appears in this case due to scaling the domain and the partial derivatives to the hypercube $[-1, 1]^m$ with a weight function 2^{-m}. Using Theorem 4.1,

$$\nu_i \;=\; \alpha_i(n) + \sum_{j=n+1}^{m} \lambda_j w_{i,j}^2 \;\leq\; \alpha_i(n) + \lambda_{n+1} \sum_{j=n+1}^{m} w_{i,j}^2 \;\leq\; \alpha_i(n) + \lambda_{n+1}. \quad (29)$$

The last inequality follows since the rows of $\boldsymbol{W}$ have norm 1. Combining (28) with (29) completes the proof.

Theorem 4.2 tells us that if λ_{n+1} is small—which often indicates an n-dimensional active subspace—and $\alpha_i(n)$ is small, then the corresponding component of the total sensitivity index will also be small. However, Sobol and Kucherenko show an example of a function in [28, Section 7] where the total sensitivity indices and the derivative-based global sensitivity measures identify different sets of variables as important. The situation is similar for the activity scores. This should not be surprising since different sensitivity metrics measure different characteristics of the function. The derivative-based metrics measure the average response to small perturbations in the inputs; the activity scores are comparable when the function admits an active subspace. In contrast, the total sensitivity indices measure the variance attributable to each parameter. In many nicely behaved functions derived from practical engineering models (such as the examples in section 5), all metrics are consistent. However, it is possible to construct functions where the metrics induce different rankings. The appropriate choice of metric depends on the application.

If the dimension of $\mathbf{x}$ is sufficiently small and the derivatives of f are sufficiently smooth, then one may estimate the integrals defining the matrix $\boldsymbol{C}$ in (14) using a high order numerical quadrature rule, such as tensor product Gauss-Legendre quadrature. Then the eigenvalues are computed from the numerically integrated matrix; we use this approach to compute reference values in section 5.

When these conditions are not satisfied (e.g., if the model contains more than a handful of input parameters), which is often the case in complex simulation models, one may get coarser estimates of the activity scores α_i from the eigenpairs $\hat{\Lambda}$, $\hat{\boldsymbol{W}}$ of the Monte Carlo estimate $\hat{\boldsymbol{C}}$ in (18). Our previous work studied the approximation properties of the Monte Carlo-based estimates of the eigenvalues and the subspaces [6]. Unfortunately, those

results do not directly translate to a priori error measures on the estimated activity scores. We are currently working to find a lower bound on the number M of samples in (18) such that the activity scores are accurate to within a user-specified tolerance—similar to [6, Corollary 3.3], which bounds the number of samples needed for accurate eigenvalue estimation. We conjecture that this number may be smaller than the number of samples needed for accurate Monte Carlo estimates of the derivative-based metrics ν_i from (9), but such analysis is beyond the scope of the current paper.

Since the Monte Carlo-based estimates of the eigenpairs cannot be interpreted as sums of independent random variables, we cannot compute central limit theorem-based standard errors. This limitation extends to the activity scores derived from the eigenpairs. Therefore, we use the bootstrap to estimate the standard error in the Monte Carlo-based estimates of the activity scores. The bootstrap algorithm is described in detail by Efron and Tibshirani [26, Algorithm 6.1]. We denote the bootstrap standard error of the activity score α_i by $\widehat{se}_{\alpha_i}$. In section 5, we study the bootstrap standard error's behavior as the number M of samples increases in the Monte Carlo-based estimates.

5. Numerical experiments

The following numerical experiments support the claim that the proposed sensitivity metrics—namely, the first eigenvector from Section 4.1 and the activity scores from Section 4.2—are consistent with standard sensitivity metrics. We test two models from Bingham's Virtual Library of Simulation Experiments [31]: (i) a nonlinear, algebraic model of a cylindrical piston with $m = 7$ physical input parameters and (ii) a nonlinear, algebraic model of a transformerless push-pull circuit with $m = 6$ physical input parameters. We chose these two models for two reasons. First, the algebraic expressions allow calculation of analytic gradients that are straightforward to implement, thus enabling the gradient-based metrics. Second, the models' fast numerical evaluation and low parameter space dimension allow us to compute accurate Gauss quadrature-based reference values for all quantities. Moreover, the fast evaluation enables a thorough error and convergence study for the Monte Carlo estimates of all metrics. In total, the numerical experiments used several million evaluations of each model, which would not have been possible with a more complex, expensive simulation model. We emphasize that the goal is not to study these models but to compare the sensitivity metrics.

The parameter space in each model is a hyperrectangle defined by independent ranges on each parameter. For our studies, we normalize the parameter space to the $[-1, 1]$ hypercube and scale the partial derivatives appropriately. Consequently, the sensitivity metrics measure sensitivity for the normalized model. We note that the parameter ranges are part of the definition of the model; the sensitivity metrics and active subspace computations depend on the given ranges.

For each model, we compute the active subspaces by (i) estimating the matrix $\boldsymbol{C}$ from (14) using tensor product Gauss-Legendre quadrature with 7 points per dimension and (ii) computing the eigenpairs of the quadrature-based estimate. We performed a convergence study for the quadrature-based estimates, and we found that with 7 points per dimension (i.e., 7^m total points), the quadrature-based estimates of the eigenvalues are accurate to 10 digits. We plot the eigenvalues on a logarithmic scale, and we find that the eigenvalues—in both models—decay spectrally, which is comparable to the real applications we have studied [4, 7, 12, 13]. We also plot the components of the first two eigenvectors, which define the first two active variables. Using the eigenvectors, we plot the one- and two-dimensional summary plots for 500 random samples of the input/output pairs. The summary plots show the relationship between the first two active variables (i.e., the two most important linear combinations of the model inputs, where the weights of the linear combinations are the first two eigenvectors of $\boldsymbol{C}$) and the output quantity of interest. If a simple univariate relationship is present in the one-dimensional summary plot, then the output can be modeled with the low-dimensional approximation from (17). The univariate function g can be fit from the samples of the active variable / output pairs; see [5, Chapter 4] for more details. If there is no apparent univariate relationship in the one-dimensional summary plot, then the two dimensional summary plot can be used similarly. If there is no apparent relationship in the two-dimensional summary plot and the eigenvalues suggest the active subspace of dimension greater than 2 is most dominant, then the summary plots are not especially helpful in revealing the low-dimensional structure in the map from inputs to outputs. Summary plots were developed in the context of sufficient dimension reduction for statistical regression. See [32] for details in the regression context, and see [5, Chapter 1] for details in the active subspaces context. We discuss the conclusions one may draw about the active subspace from these plots before studying the sensitivity metrics.

For each of the two models, we examine the following five sensitivity

metrics:

1. the Sobol' total sensitivity index τ_i from Section 2.1,
2. the derivative-based global sensitivity measures ν_i from Section 2.2,
3. the linear model coefficients β_i from Section 2.3,
4. the first eigenvector components $w_{i,1}$ from Section 4.1,
5. and activity scores $\alpha_i(n)$, where n is the dimension of the active subspace, from Section 4.2.

The subscript i gives the index of the corresponding parameter. For each sensitivity metric, we compute

- a tensor product Gauss-Legendre quadrature-based reference value with 7 points per dimension (i.e., 7^m total points)[1],
- a Monte Carlo estimate with M samples, where M ranges between 50 and 50000,
- a relative error between the Monte Carlo estimate and the quadrature-based reference value,
- and the (bootstrap) standard error for each Monte Carlo estimate.

The number M of evaluations of $f(\mathbf{x})$ is a measure of the sensitivity metrics' cost. For the total sensitivity indices, we choose M' such that $M = (m + 1)M'$, and we use M' as the number of samples in $\mathbf{f}_{\boldsymbol{A}}$ and each $\mathbf{f}_{\boldsymbol{B}_i}$ from Section 2.1. This keeps the comparison fair, but the number of samples used to compute the integrals is effectively less than for the other metrics. Consequently, we expect the error and standard error to be larger for the total sensitivity indices. For the derivative-based metrics, we count each gradient sample as one evaluation, which is the case for simulation models with algorithmic differentiation capabilities.

We scale the error and standard error in the following manner. Let γ_i be the reference value for one of the five sensitivity metrics listed above, and let $\hat{\gamma}_i$ be its Monte Carlo estimate. We compute the relative error in the Monte Carlo estimate as

$$\frac{|\gamma_i - \hat{\gamma}_i|}{\max\limits_{1 \le j \le m} |\gamma_j|}. \tag{30}$$

[1]With 7 points per dimension, the reference values contain 10 digits of accuracy in all cases.

Scaling error by the magnitude of the largest-in-magnitude metric avoids issues of dividing by very small sensitivity metrics. Since these errors are random due to the random sampling, we average the errors, as a function of the number M of samples, over 10 independent trials. We divide the (bootstrap) standard errors by the magnitude of the largest-in-magnitude Monte Carlo estimate of the sensitivity metric; in other words, we do not use the reference values when computing the standard errors as would be done in practice without reference values. We average the standard error estimates over 10 independent trials to reduce the effects of randomness. For the active subspace-based activity scores, we should how the values change as a function of the dimension n of the active subspace. Recall that Theorem 4.1 shows that the activity scores become the derivative-based global sensitivity metrics when the dimension n of the active subspace reaches the dimension m of the parameter space.

Finally, we show that the rankings on the input variables' importance induced by the five sensitivity metrics are consistent for both models. We emphasize that this is not case in general. Since the various metrics are derived from different features of output (e.g., averaged derivatives versus variances), it is possible to devise functions where the metrics are not consistent; see, e.g., [28, Section 7]. The purpose of comparing rankings is to show that the metrics derived from the active subspace are consistent with more standard global sensitivity metrics in an engineering inspired model, which supports the analysis from Section 4.3.

The Matlab scripts for all computations and plots are available at `https://bitbucket.org/paulcon/global-sensitivity-metrics-from-active-subspaces`. All computations were run in Matlab 2015b on a MacBook Pro with a dual core processor and 16GB of memory.

5.1. Piston model

The first model is an algebraic expression for a cylindrical piston. Details of the model and associated Matlab code are available at `http://www.sfu.ca/~ssurjano/piston.html`. The piston model appears in [33, 34] as a test model for statistical screening. The quantity of interest, t, is the time in seconds it takes the piston to complete one cycle. This time depends on

$m = 7$ physical parameters according to the following nonlinear expressions:

$$t = 2\pi\sqrt{\frac{M}{k + S^2\frac{P_0V_0}{T_0}\frac{T_a}{V^2}}}, \tag{31}$$

$$V = \frac{S}{2k}\left(\sqrt{A^2 + 4k\frac{P_0V_0}{T_0}T_a} - A\right), \tag{32}$$

$$A = P_0S + 19.62M - \frac{kV_0}{S}. \tag{33}$$

The input parameters' descriptions, ranges, and units are in Table 1.

Table 1: Input parameters' descriptions, ranges, and units for the piston model (31).

Parameter	Notation	Min	Max	Units
piston weight	M	30	60	kg
piston surface area	S	0.005	0.020	m^2
initial gas volume	V_0	0.002	0.010	m^3
spring coefficient	k	1000	5000	N/m
atmospheric pressure	P_0	90000	110000	N/m^2
ambient temperature	T_a	290	296	K
filling gas temperature	T_0	340	360	K

Figure 1a shows the 7 eigenvalues, on a logarithmic scale, from the quadrature-based estimate of $\boldsymbol{C}$ from (14). The order-of-magnitude gaps between the eigenvalues suggest that an active subspace exists for n from 1 to 6. Figure 1b shows the components of the first two eigenvectors of $\boldsymbol{C}$, which are used to produce the one- and two-dimensional summary plots in Figures 1c and 1d, respectively. The one-dimensional summary plot shows a gross monotonic trend in the output as a function of the first active variable. This trend might be useful in certain applications, such as determining a good starting value for a numerical optimization routine. See [9] for an example of exploiting such a relationship for estimating a range of outputs in a multiphysics scramjet model. However, there is also significant spread around a perceived univariate relationship, so an approximation of the form (17) with one eigenvector may not be useful for all response surface applications. This spread is mostly likely due to changes in the output as the input is perturbed orthogonally to the first eigenvector of $\boldsymbol{C}$. The two-dimensional

summary plot captures more of the input/output relationship, where some curvature is apparent in the two-dimensional level sets. Thus, an approximation of the form (17) with two eigenvectors may be more useful than with only one eigenvector. See [7] for an example of building a response surface that exploits the two-dimensional active subspace for aerospace shape design. This sort of qualitative reasoning is typical when the engineer is determining how to exploit the active subspace for a particular application.

Table 2 shows the reference values for each sensitivity metric to four digits. Since the eigenvalue gap is largest between the first and second eigenvalues, we choose $n = 1$ when computing the activity scores to compare to the other sensitivity metrics. These are the values used to compute the error of the Monte Carlo estimates.

Table 2: Reference values for piston model's sensitivity metrics computed with tensor product Gauss-Legendre quadrature. The metrics compared are the Sobol' total sensitivity index (τ_i, (5)), the derivative-based global sensitivity measure (ν_i, (9)), the linear model coefficients (β_i, (13)), the first eigenvector components ($w_{i,1}$, Section 4.1), and the activity scores (α_i, (21)).

Parameter	τ_i	ν_i	β_i	$w_{i,1}$	$\alpha_i(1)$
M	0.0509	0.0032	0.1963	0.1604	0.0018
S	0.5994	0.0449	-0.7345	-0.7936	0.0437
V_0	0.3528	0.0265	0.5567	0.5768	0.0231
k	0.0669	0.0040	-0.1424	-0.1035	0.0007
P_0	0.0013	0.0001	-0.0352	-0.0305	0.0001
T_a	0.0000	0.0000	0.0018	0.0015	0.0000
T_0	0.0001	0.0000	-0.0051	-0.0042	0.0000

Figure 2 shows relative error in the Monte Carlo estimates of the sensitivity metrics, averaged over 10 independent trials. All errors decrease at the expected $M^{-1/2}$ convergence rate. The errors for the total sensitivity indices are larger due to how we allocate the M model evaluations as described above. Across all metrics, the errors in the large metrics are relatively larger than the errors in the small metrics—with the notable exception of the linear model coefficients, which have roughly the same error. This is likely due to our choice for normalization in (30).

Figure 3 shows the (bootstrap) standard errors for the Monte Carlo estimates. The bootstrap standard errors are generally lower than the relative

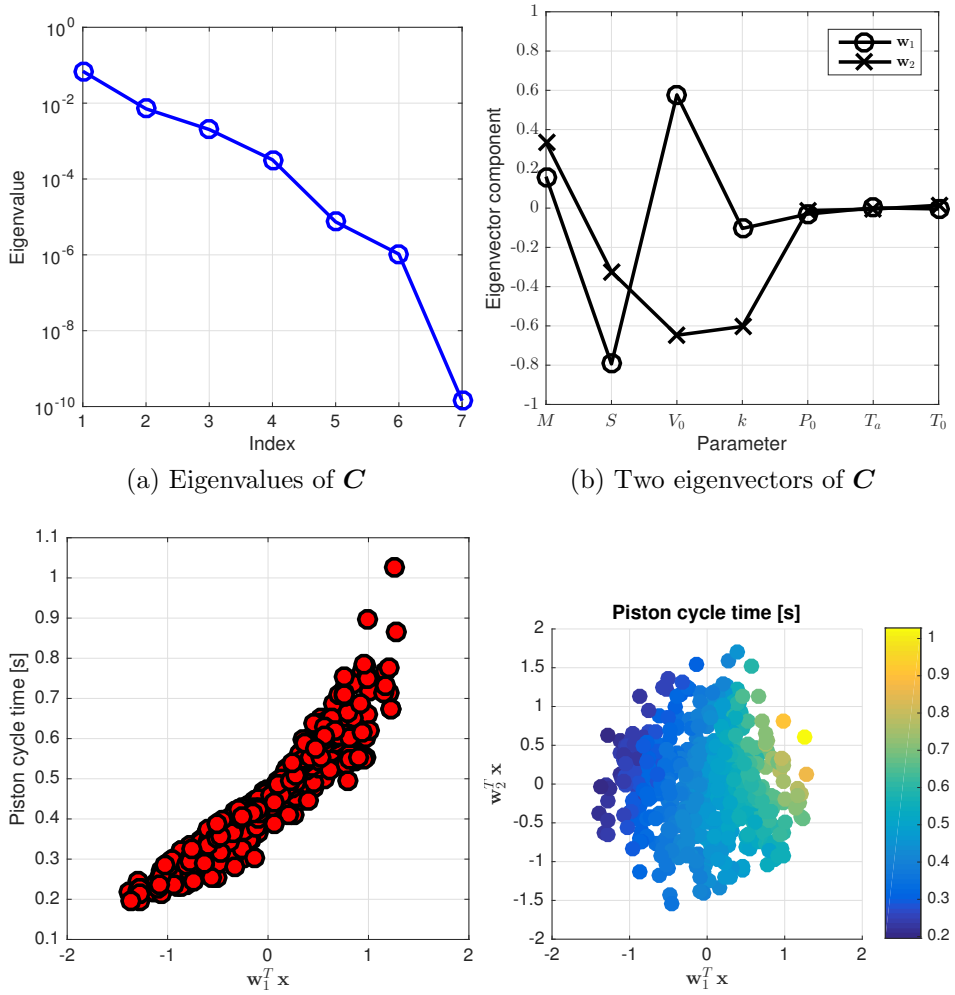

(a) Eigenvalues of $\boldsymbol{C}$

(b) Two eigenvectors of $\boldsymbol{C}$

(c) 1D summary plot

(d) 2D summary plot

Figure 1: The eigenvalues in the top left subfigure provide evidence of active subspaces of dimension 1 to 6 in the piston cycle time as a function of its 7 physical input parameters from Table 1. Components of the first two eigenvectors are in the top right. The one- and two-dimensional summary plots in the second row elucidate the relationship between the first two active variables and the piston cycle time.

errors from Figure 2. However, all standard errors decay at the expected rate of $M^{-1/2}$.

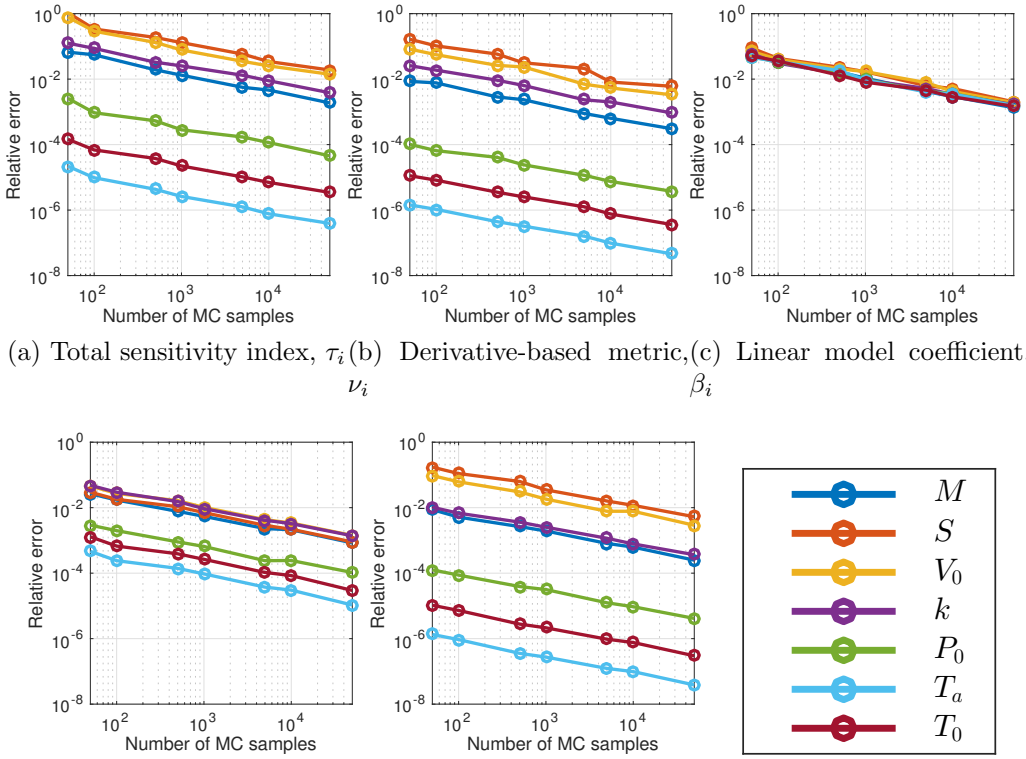

(a) Total sensitivity index, τ_i (b) Derivative-based metric, ν_i (c) Linear model coefficient, β_i

(d) First eigenvector, $w_{i,1}$ (e) Activity score, $\alpha_i(1)$ (f) Legend

Figure 2: The relative error in the Monte Carlo estimates of the sensitivity metrics for the piston model (31) as a function of the number M of samples, computed using the reference values in Table 2 and the formula in (30).

5.2. Circuit model

The second model is an algebraic expression for a transformerless push-pull circuit. Details of the model and associated Matlab code are available at `http://www.sfu.ca/~ssurjano/otlcircuit.html`. The circuit model also appears in [33, 34] as a test model for statistical screening. The quantity of interest, V, is the midpoint voltage of the circuit. This voltage depends on

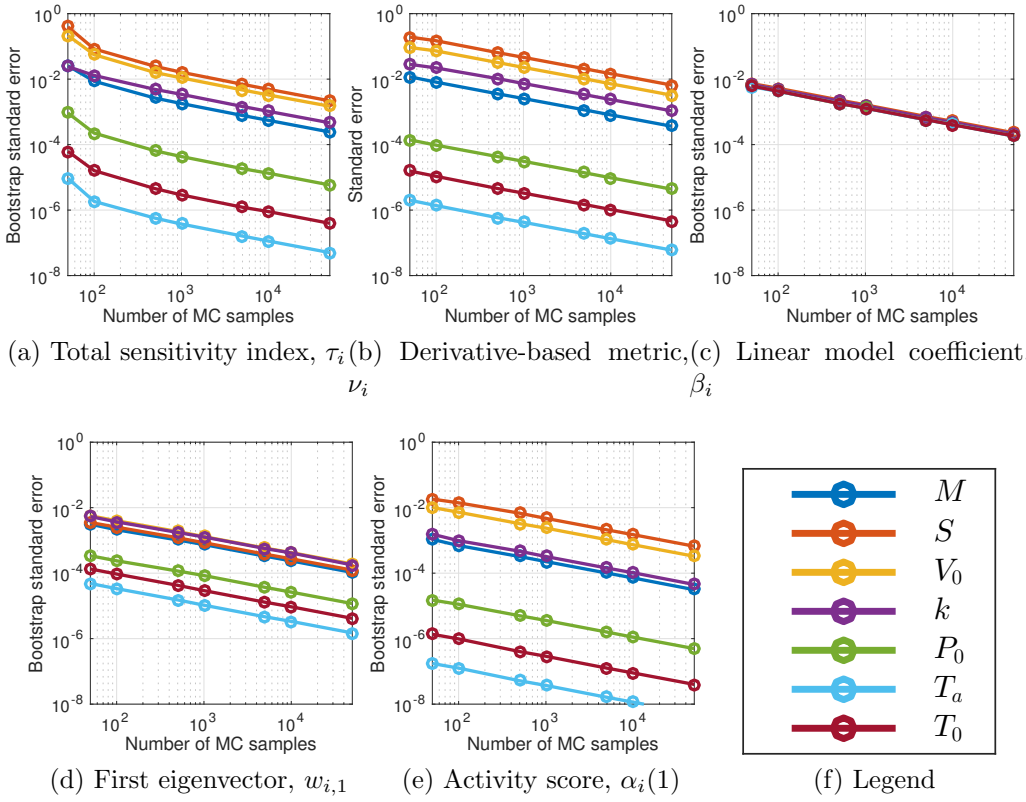

(a) Total sensitivity index, τ_i (b) Derivative-based metric, ν_i (c) Linear model coefficient, β_i

(d) First eigenvector, $w_{i,1}$ (e) Activity score, $\alpha_i(1)$ (f) Legend

Figure 3: The (bootstrap) standard errors in the Monte Carlo estimates of the sensitivity metrics for the piston model (31) as a function of the number M of samples.

$m = 6$ physical parameters according to the following nonlinear expressions:

$$V = \frac{(12R_{b2}/(R_{b1}+R_{b2})+0.74)\beta(R_{c2}+9)}{\beta(R_{c2}+9)+R_f} \tag{34}$$

$$+ \frac{11.35R_f}{\beta(R_{c2}+9)+R_f} + \frac{0.74R_f\beta(R_{c2}+9)}{(\beta(R_{c2}+9)+R_f)R_{c1}}. \tag{35}$$

The input parameters' descriptions, ranges, and units are in Table 3.

Figure 4a shows the 6 eigenvalues, on a logarithmic scale, from the quadrature-based estimate of $\boldsymbol{C}$ from (14). The order-of-magnitude gaps between the eigenvalues suggest that an active subspace exists for n from 1 to 5. Figure 4b shows the components of the first two eigenvectors of $\boldsymbol{C}$,

Table 3: Input parameters' descriptions, ranges, and units for the circuit model (34).

Parameter	Notation	Min	Max	Units
resistance b1	R_{b1}	50	150	K-Ohms
resistance b2	R_{b2}	25	70	K-Ohms
resistance f	R_f	0.5	3.0	K-Ohms
resistance c1	R_{c1}	1.2	2.5	K-Ohms
resistance c2	R_{c2}	0.25	1.20	K-Ohms
current gains	β	50	300	Amperes

which are used to produce the one- and two-dimensional summary plots in Figures 4c and 4d, respectively. The one-dimensional summary plot shows a nearly linear relationship in the output as a function of the first active variable. Compared to the piston model (see Figure 1c), the univariate relationship is tighter, and the spread away from the univariate relationship is relatively small. This suggests that, relative to the structure in the piston model, the relationship between the circuit parameters and the output voltage can be better modeled by a function—even a linear function—of the first active variable. In this case g from (17) may be a univariate linear function fit with least-squares; see [5, Chapter 4] for more details. This is confirmed by the two-dimensional summary plot, which has nearly linear contours from the 500 random samples. Again, this is the sort of qualitative reasoning an engineer may use to determine how to exploit the low-dimensional structure for a particular application.

Table 4 shows the reference values for each sensitivity metric to four digits. Similar to the piston model, since the eigenvalue gap is largest between the first and second eigenvalues, we choose $n = 1$ when computing the activity scores to compare to the other sensitivity metrics. These are the values used to compute the error of the Monte Carlo estimates.

Figure 5 shows relative errors in the Monte Carlo estimates of the sensitivity metrics, averaged over 10 independent trials. All errors decrease at the expected $M^{-1/2}$ convergence rate. The errors in the total sensitivity indices are relatively large due to how we allocate the samples. Similar to the piston model, the errors in the large metrics are relatively larger than the errors in the small metrics—except for the linear model coefficients, which have roughly the same error.

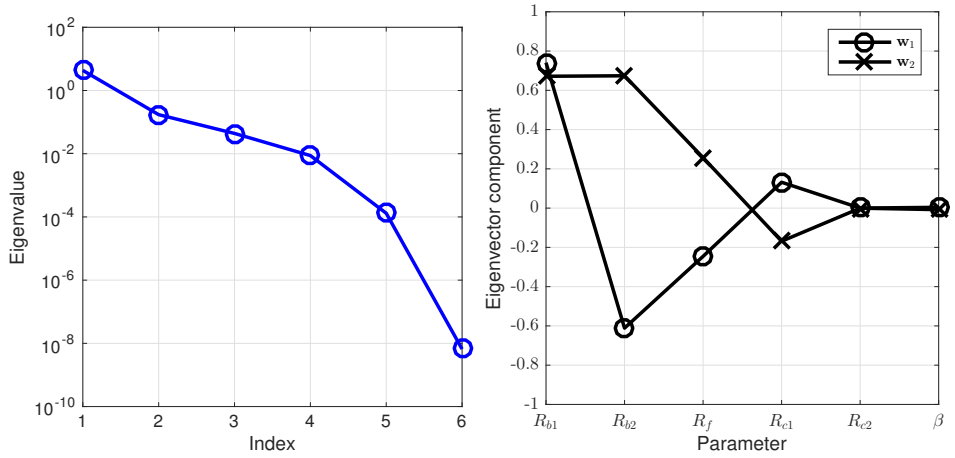

(a) Eigenvalues of $\boldsymbol{C}$ (b) Two eigenvectors of $\boldsymbol{C}$

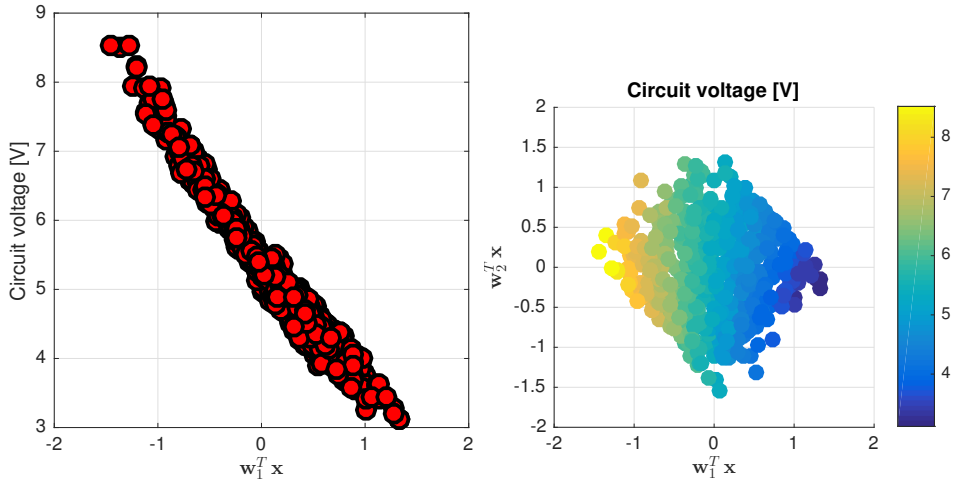

(c) 1D summary plot (d) 2D summary plot

Figure 4: The eigenvalues in the top left subfigure provide evidence of active subspaces of dimension 1 to 5 in the circuit voltage as a function of its 6 physical input parameters from Table 3. Components of the first two eigenvectors are in the top right. The one- and two-dimensional summary plots in the second row elucidate the relationship between the first two active variables and the circuit voltage.

Figure 6 shows the (bootstrap) standard errors for the Monte Carlo estimates. Again, the bootstrap standard errors are generally lower than the

Table 4: Reference values for circuit model's sensitivity metrics. The notation is similar to Table 2.

Parameter	τ_i	ν_i	β_i	$w_{i,1}$	$\alpha_i(1)$
R_{b1}	0.5001	2.4555	-0.6925	-0.7407	2.3778
R_{b2}	0.4117	1.7039	0.6358	0.6112	1.6190
R_f	0.0740	0.2902	0.2662	0.2458	0.2617
R_{c1}	0.0218	0.1090	-0.1341	-0.1318	0.0752
R_{c2}	0.0000	0.0000	-0.0002	-0.0002	0.0000
β	0.0000	0.0002	-0.0032	-0.0039	0.0001

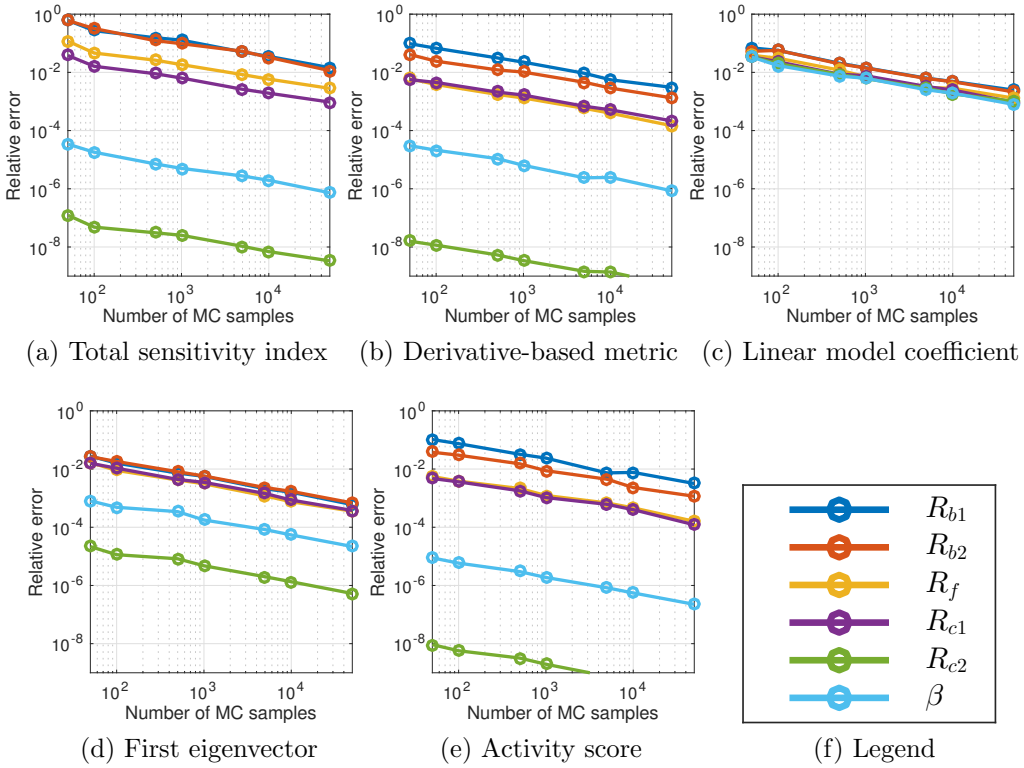

(a) Total sensitivity index (b) Derivative-based metric (c) Linear model coefficient

(d) First eigenvector (e) Activity score (f) Legend

Figure 5: The relative error in the Monte Carlo estimates of the sensitivity metrics for the circuit model (34) as a function of the number M of samples, averaged over 10 independent trials and computed using the error definition (30) with reference values from Table 3.

relative errors from Figure 2. And all standard errors decay at the expected rate of $M^{-1/2}$.

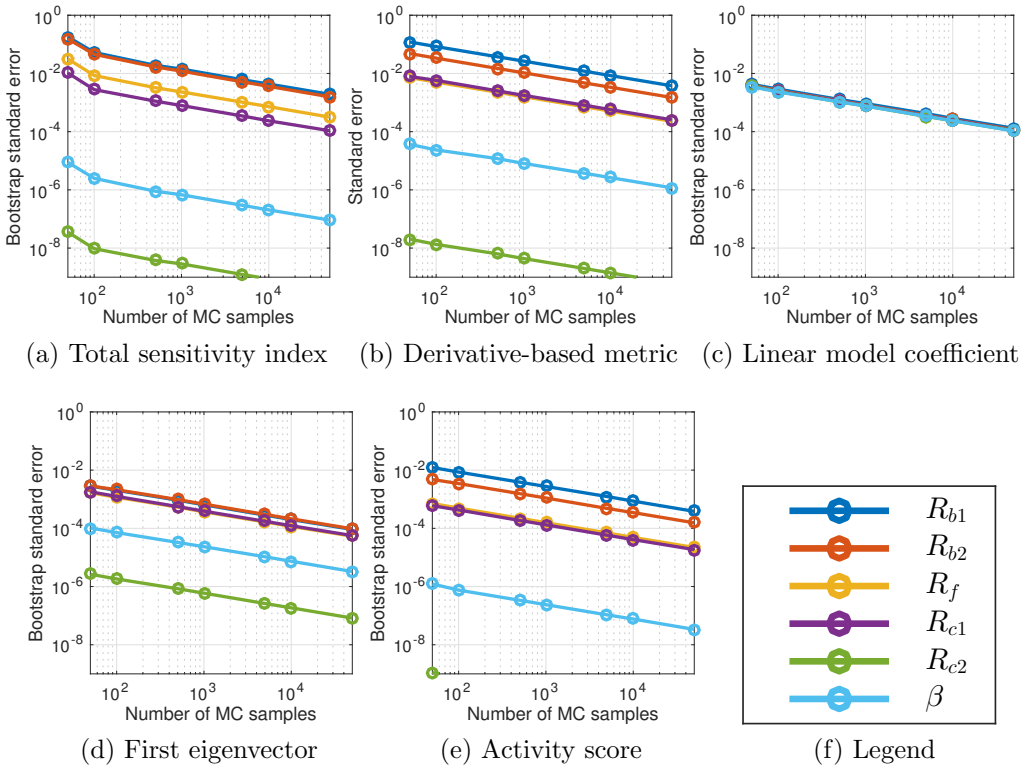

(a) Total sensitivity index (b) Derivative-based metric (c) Linear model coefficient

(d) First eigenvector (e) Activity score (f) Legend

Figure 6: The (bootstrap) standard errors for the Monte Carlo estimates as a function of the number M of samples averaged over 10 independent trials.

5.3. Comparing rankings across metrics

Figure 7 shows the activity scores as a function of the active subspace dimension n for each model. The rankings are consistent as n increases with one exception: in the piston model, the rankings of parameters M and k swap as the dimension n increases from 1 to 2. The activity scores converge quickly, which is due to the spectral decay of the eigenvalues weighting the squared eigenvector components in (21). Note that the lines representing the activity scores for parameters P_0 and T_a in the piston model are behind the line for T_0, and all are near 0. Similarly, the lines representing the activity scores for parameter R_{c2} are behind the line for β in the circuit model.

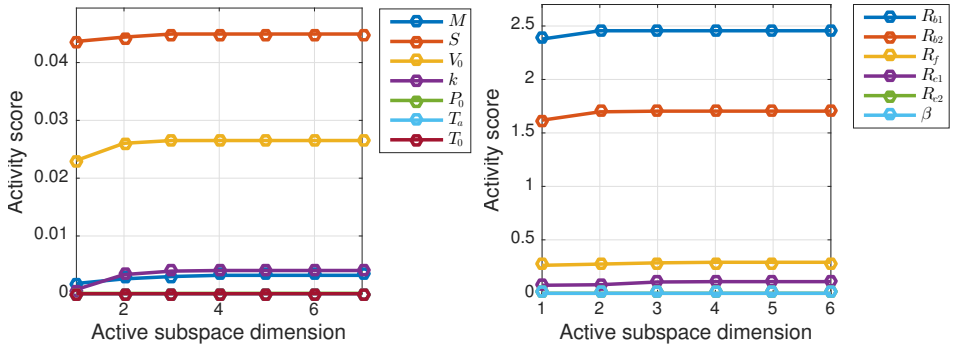

(a) Piston model (b) Circuit model

Figure 7: Activity scores as a function of the active subspace dimension n. Recall that when $n = m$, the activity scores are the same as the derivative-based global sensitivity measures, as shown in Theorem 4.1.

To visually compare rankings, we plot normalized versions of the reference values of the sensitivity metrics for each model. The normalization is as follows. For a sensitivity metric γ_i, we plot

$$\frac{|\gamma_i|}{\left(\sum_{i=1}^{m} \gamma_i^2\right)^{1/2}}, \tag{36}$$

which (i) removes the signs from the linear model coefficients and the first eigenvector's components and (ii) scales the vector of numerical values to have norm 1. Figure 8 shows the normalized versions of the sensitivity metrics for both models. The parameters' rankings are consistent across the metrics for these two models. This supports our claim that the active subspace-based sensitivity metrics can be used like standard sensitivity metrics to measure the parameters' importance—at least, for these two models. And this complements the analysis from Section 4.3.

6. Summary and remarks

We propose two global sensitivity metrics derived from the eigenvectors and eigenvalues computed while testing an engineering model for an active subspace using its gradients. The two metrics are (i) the components of the first eigenvector and (ii) the *activity scores*, which are the squared eigenvector components, linearly combined with the eigenvalues. We mathemati-

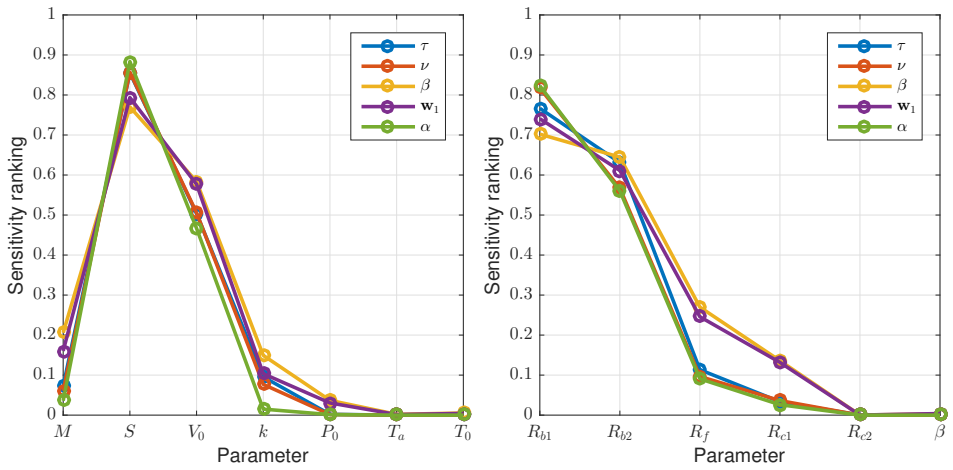

(a) Piston model (b) Circuit model

Figure 8: Normalized sensitivity metrics allow visual comparison of importance rankings from the various sensitivity metrics for the two models. All metrics rank the parameters consistently for these two models.

cally relate the activity scores to the Sobol' total sensitivity indices and the derivative-based global sensitivity measures. We discuss Gauss quadrature methods and Monte Carlo methods to estimate the new sensitivity metrics. And we demonstrate the new metrics on two algebraic engineering models. The models are simple enough to compute accurate quadrature-based reference values and study the behavior of the Monte Carlo estimates and their (bootstrap) standard errors estimates. For the two test models, all metrics identify the same important and unimportant variables, which supports our claim that the active subspace-based metrics behave similarly to the standard sensitivity metrics. While it is possible to construct models where the metrics rank parameters differently, since the metrics measure different characteristics of the model, we expect that this consistency will occur across many models in engineering practice.

Our thesis is that the sensitivity metrics we derive from active subspaces behave similarly to the standard sensitivity metrics. We do not claim that the proposed metrics are computationally advantageous. In fact, the numerical experiments from Section 5 suggest that the error in the Monte Carlo estimates of each metric are comparable for a fixed number M of samples—particularly when comparing the activity scores to the derivative-based sen-

sitivity measures. Similarly, the (bootstrap) standard errors are comparable across metrics, which suggests that the Monte Carlo estimates' variances are comparable for the same number of samples. Future work will analyze the error in the activity scores' Monte Carlo estimates as a function of the number of samples. Such analysis is beyond the scope of the current paper.

Similar to the derivative-based global sensitivity measures, estimating the active subspace requires access to gradients. If gradients are not available and finite differences are not feasible, then one may seek to estimate active subspaces without gradients. Some initial work has been done along these lines [35, 36]; we plan to explore how such gradient-free approaches lead to sensitivity metrics.

Acknowledgments

This work is supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC-0011077.

References

[1] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, S. Tarantola, Global Sensitivity Analysis: The Primer, John Wiley & Sons, New York, 2008.

[2] I. Sobol, Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, Mathematics and Computers in Simulation 55 (2001) 271–280, The Second IMACS Seminar on Monte Carlo Methods.
URL http://dx.doi.org/10.1016/S0378-4754(00)00270-6

[3] S. Kucherenko, M. Rodriguez-Fernandez, C. Pantelides, N. Shah, Monte Carlo evaluation of derivative-based global sensitivity measures, Reliability Engineering & System Safety 94 (7) (2009) 1135–1148, Special Issue on Sensitivity Analysis.
URL http://dx.doi.org/10.1016/j.ress.2008.05.006

[4] P. Constantine, E. Dow, Q. Wang, Active subspace methods in theory and practice: Applications to kriging surfaces, SIAM Journal on Scientific Computing 36 (4) (2014) A1500–A1524.
URL http://dx.doi.org/10.1137/130916138

[5] P. G. Constantine, Active Subspaces: Emerging Ideas in Dimension Reduction for Parameter Studies, SIAM, Philadelphia, 2015.
URL http://dx.doi.org/10.1137/1.9781611973860

[6] P. Constantine, D. Gleich, Computing active subspaces with Monte Carlo, arXiv preprint arXiv:1408.0545v2.
URL http://arxiv.org/abs/1408.0545

[7] T. W. Lukaczyk, P. Constantine, F. Palacios, J. J. Alonso, Active subspaces for shape optimization, 10th AIAA Multidisciplinary Design Optimization Conference, National Harbor, MDAIAA 2014-1171.
URL http://dx.doi.org/10.2514/6.2014-1171

[8] C. Othmer, T. W. Lukaczyk, P. Constantine, J. J. Alonso, On active subspaces in car aerodynamics, 17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Washington, D.C.AIAA 2016-4294.
URL http://dx.doi.org/10.2514/6.2016-4294

[9] P. Constantine, M. Emory, J. Larsson, G. Iaccarino, Exploiting active subspaces to quantify uncertainty in the numerical simulation of the HyShot II scramjet, Journal of Computational Physics 302 (2015) 1–20.
URL http://dx.doi.org/10.1016/j.jcp.2015.09.001

[10] J. L. Jefferson, J. M. Gilbert, P. G. Constantine, R. M. Maxwell, Active subspaces for sensitivity analysis and dimension reduction of an integrated hydrologic model, Computers & Geosciences 83 (2015) 127–138.
URL http://dx.doi.org/10.1016/j.cageo.2015.07.001

[11] J. M. Gilbert, J. L. Jefferson, P. G. Constantine, R. M. Maxwell, Global spatial sensitivity of runoff to subsurface permeability using the active subspace method, Advances in Water Resources 92 (2016) 30–42.
URL http://dx.doi.org/10.1016/j.advwatres.2016.03.020

[12] P. G. Constantine, B. Zaharatos, M. Campanelli, Discovering an active subspace in a single-diode solar cell model, Statistical Analysis and Data Mining: The ASA Data Science Journal 8 (5-6) (2015) 264–273.
URL http://dx.doi.org/10.1002/sam.11281

[13] M. Bauerheim, A. Ndiaye, P. Constantine, S. Moreau, F. Nicoud, Symmetry breaking of azimuthal thermoacoustic modes: the UQ perspective,

Journal of Fluid Mechanics 789 (2016) 534–566.
URL http://dx.doi.org/10.1017/jfm.2015.730

[14] P. G. Constantine, A. Doostan, Time-dependent global sensitivity analysis with active subspaces for a lithium ion battery model, arXiv preprint arXiv:1606.08770.
URL https://arxiv.org/abs/1606.08770

[15] P. Diaz, P. Constantine, K. Kalmbach, E. Jones, S. Pankavich, A modified SEIR model for the spread of Ebola in western Africa and metrics for resource allocation, arXiv preprint arXiv:1603.04955.
URL http://arxiv.org/abs/1603.04955

[16] P. G. Constantine, Z. del Rosario, G. Iaccarino, Many physical laws are ridge functions, arXiv preprint arXiv:1605.07974.
URL http://arxiv.org/abs/1605.07974

[17] P. G. Constantine, A. Eftekhari, R. Ward, A near-stationary subspace for ridge approximation, arXiv preprint arXiv:1606.01929.
URL http://arxiv.org/abs/1606.01929

[18] A. Owen, Variance components and generalized Sobol' indices, SIAM/ASA Journal on Uncertainty Quantification 1 (1) (2013) 19–41.
URL http://dx.doi.org/10.1137/120876782

[19] R. Liu, A. B. Owen, Estimating mean dimensionality of analysis of variance decompositions, Journal of the American Statistical Association 101 (474) (2006) 712–721.
URL http://www.jstor.org/stable/27590729

[20] I. Sobol', E. Myshetskaya, Monte Carlo estimators for small sensitivity indices, Monte Carlo Methods & Applications 13 (5/6) (2008) 455–465.
URL http://dx.doi.org/10.1515/mcma.2007.023

[21] B. Sudret, Global sensitivity analysis using polynomial chaos expansions, Reliability Engineering & System Safety 93 (7) (2008) 964–979, bayesian Networks in Dependability.
URL http://dx.doi.org/10.1016/j.ress.2007.04.002

[22] T. Crestaux, O. L. Maître, J.-M. Martinez, Polynomial chaos expansion for sensitivity analysis, Reliability Engineering & System Safety 94 (7)

(2009) 1161–1172.
URL http://dx.doi.org/10.1016/j.ress.2008.10.008

[23] A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, S. Tarantola, Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index, Computer Physics Communications 181 (2) (2010) 259–270.
URL http://dx.doi.org/10.1016/j.cpc.2009.09.018

[24] V. G. Weirs, J. R. Kamm, L. P. Swiler, S. Tarantola, M. Ratto, B. M. Adams, W. J. Rider, M. S. Eldred, Sensitivity analysis techniques applied to a system of hyperbolic conservation laws, Reliability Engineering & System Safety 107 (2012) 157–170, SAMO 2010.
URL http://dx.doi.org/10.1016/j.ress.2011.12.0087

[25] M. J. Jansen, Analysis of variance designs for model output, Computer Physics Communications 117 (1) (1999) 35–43.
URL http://dx.doi.org/10.1016/S0010-4655(98)00154-4

[26] B. Efron, R. J. Tibshirani, An Introduction to the Bootstrap, Chapman and Hall/CRC, Boca Raton, 1993.

[27] M. D. Morris, Factorial sampling plans for preliminary computational experiments, Technometrics 33 (2) (1991) 161–174.
URL http://dx.doi.org/10.1080/00401706.1991.10484804

[28] I. Sobol', S. Kucherenko, Derivative based global sensitivity measures and their link with global sensitivity indices, Mathematics and Computers in Simulation 79 (10) (2009) 3009–3017.
URL http://dx.doi.org/10.1016/j.matcom.2009.01.023

[29] D. Xiu, G. E. Karniadakis, The Wiener-Askey polynomial chaos for stochastic differential equations, SIAM Journal on Scientific Computing 24 (2) (2002) 619–644.
URL http://dx.doi.org/10.1137/S1064827501387826

[30] P. G. Constantine, M. S. Eldred, E. T. Phipps, Sparse pseudospectral approximation method, Computer Methods in Applied Mechanics and Engineering 229–232 (2012) 1–12.
URL http://dx.doi.org/10.1016/j.cma.2012.03.019

[31] Virtual library of simulation experiments: Test functions and datasets, http://www.sfu.ca/~ssurjano/index.html, accessed: 10/13/2015.

[32] R. D. Cook, Regression Graphics: Ideas for Studying Regressions through Graphics, John Wiley & Sons, Hoboken, 1998.

[33] E. N. Ben-Ari, D. M. Steinberg, Modeling data from computer experiments: an empirical comparison of kriging with MARS and projection pursuit regression, Quality Engineering 19 (4) (2007) 327–338.
URL http://dx.doi.org/10.1080/08982110701580930

[34] H. Moon, Design and analysis of computer experiments for screening input variables, Ph.D. thesis, The Ohio State University (2010).

[35] P. G. Constantine, A. Eftekhari, M. B. Wakin, Computing active subspaces efficiently with gradient sketching, in: 2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2015, pp. 353–356.
URL http://dx.doi.org/10.1109/CAMSAP.2015.7383809

[36] W. Li, G. Lin, B. Li, Inverse regression-based uncertainty quantification algorithms for high-dimensional models: Theory and practice, Journal of Computational Physics 321 (2016) 259–278.
URL http://dx.doi.org/10.1016/j.jcp.2016.05.040